Guidelines for Program/Project Responsibilities for Safety and Mission Success

Risk Management: Beginning in the formulation phase, program/project managers will continuously use risk management in all decisionmaking to increase the likelihood of programmatic and technical success. Risk is anything that threatens mission success, including safety, cost, schedule, and technical risks. Program/project decisions, including decisions on the depth of safety and mission success analysis, must be made based on an orderly risk management process. Risk management is especially important if we expect to be safer and more successful with faster, better, and cheaper projects. (See appendix 1 for the text of the Administrator's discussion of “Risk Management” at the February 9, 2000, Senior Staff and Center Directors’ meeting.)

Risk Management Technologies: There exists a wide range of proven tools available to identify and analyze risks to safety and mission success. Beginning in the formulation phase, program/project managers will constantly seek to identify and analyze what could go wrong with their programs/projects so that mitigation efforts can be identified and applied. Several proven tools for this purpose include Failure Modes and Effects Analysis, Fault Tree Analysis, and Probabilistic Risk Assessments (and there are many others including preliminary, subsystem, system, and operational hazard analysis; sneak circuit analysis, common cause failure analysis, Ishakawa or fishbone diagrams, etc.; NASA Reference Publication 1358, “System Engineering ‘Tool Box’ for Design-Oriented Engineers” briefly describes many of the available tools).

- **Failure Modes and Effects Analysis (FMEA).** The FMEA is a bottom-up analysis of component-level failures and their effects on higher-level systems; it is usually performed to identify critical hardware items. FMEA can help verify that all safety-critical hardware has been identified and addressed in hazard analyses. All credible failure modes and their resultant effects at the component and system levels are identified and documented. The analysis follows a well-defined sequence of steps that encompasses: (1) failure modes; (2) failure effects; (3) causes; (4) detectability; (5) corrective or preventative actions; and (6) acceptance rationale.

- **Fault Tree Analysis (FTA).** FTA is a top-down analysis used to evaluate specific undesired events. It is a deductive logic tree linking a top event to the combinations of sub-events that could cause it. It is primarily used for the qualitative study of hazardous events in systems, subsystems, components, or operations. FTA can verify that a FMEA has identified single failure points. It also can be used for quantitatively evaluating the probability of the top event when data are available for the contributing sub-events. (See appendix 2 for the text of the Administrator's discussion of “Fault Tree Analysis” at the January 20, 2000, Senior Staff and Center Directors’ meeting.)

- **Probabilistic Risk Assessment (PRA).** PRA provides a means for expressing quantitatively our state of knowledge about the risk of failure. It is an analysis of the probability (or frequency) of occurrence of a top-level undesired event, including an assessment and display of our degree of uncertainty surrounding the probability. PRA is based on comprehensive systems analysis and is repeated periodically as the design matures and new data become available. PRA can be used to support strategic decisionmaking such as in answering the question "What is the probability of losing the multi-billion dollar International Space Station (ISS) during its assembly?" (This question was asked by the NASA Advisory Committee; a PRA of the ISS is well underway.) For systems under development, PRA provides a basis for tradeoffs among safety, reliability, cost, performance, and other resources. For mature systems, it can be used for decisionmaking on risk acceptability, and, when risk is considered to be too high, choosing among options for risk reduction. For example, NASA is using PRA to assist in decisionmaking on Space Shuttle upgrades. It may also be used to track risk levels throughout the life cycle of a program/project.

While FMEA's, FTA's, PRA's and other safety and mission success methods are best applied early, they can provide useful results when applied at virtually any time during the program/project life cycle.

Design for Safety (and Mission Success): Program/project managers will begin early in the formulation phase to manage risk and identify and analyze specific hardware, software, and human failure modes—at the very time when it is most effective and efficient to design for safety. We will design to reduce hazard effects or eliminate hazards and for improved reliability. We will plan systems integration and test processes early so that, when the time comes, we will be able to verify with confidence that we have done a good job of designing for safety. We will use a total systems approach so that we do not neglect the potentially negative effects of subsystem and component interactions. We will exploit state-of-the-art and emerging technology, such as intelligent systems, advanced analysis tools, and the Intelligent Synthesis Environment, to find the problems that humans might otherwise miss and to produce "smart," robust designs. We will monitor risk continuously, watch for problems that were not anticipated, and use artificial intelligence techniques to spot patterns and trends. We will be smart, skeptical, vigilant, and design for safety and mission success. (See appendix 3 for an outline of the Administrator's discussion of "Design for Safety" at the February 28, 2000, Senior Staff and Center Directors' meeting.)

Program Reviews: If we all did a perfect job of managing risk, the safety and mission success aspects of program reviews could become unnecessary. But, invariably, we humans are not perfect, so reviews will continue to serve the Agency as a safety net. Program reviews from peer reviews to Headquarters Program Management Council reviews must be done well, and they must be critical in their approach. They must ensure and validate that we have done the best possible job in all areas, particularly in those that relate to safety and mission success. To explore failure causes and expose program weaknesses at these reviews, open dialog should address the following general areas:

- Minimum mission success criteria keyed to potential mission failure events;
- Identification of credible causes of mission failure;
- Identification of critical hardware and software items, the failure of which could lead to mission failure;
- Cross-check of mission failure causes against critical items;
- Probability of:
  1. Death or serious injury to members of the public;
  2. Loss of astronaut or test pilot crew;
  3. Lost-time injury to members of the NASA workforce;
  4. Loss of high-value equipment or property;
  5. Mission success.
- Risk mitigation plans.

Appendices:

A. Risk Management
B. Fault Tree Analysis
C. Design for Safety

Risk Management

Risk—safety, technical, cost, schedule, and other types can threaten mission success. To reduce risk, we need to manage our projects systematically, especially if we expect to be successful with faster, better, cheaper projects. Risk Management is not high tech and it is not complicated. The Risk Management process efficiently identifies, analyzes, plans, tracks, controls, communicates, and documents risk to increase the likelihood of achieving program/project goals.

Every project should have a prioritized list of its risks at any point in the life cycle, along with the programmatic impacts. The list should indicate which risks have the highest probability, which have the highest consequences, and which need to be worked now. It means that all members of the project team should have access to the risk list so that everyone knows what the risks are. It means that the project team members are responsible for the risks. The team should work to reduce or eliminate the risks that exist and develop contingency plans, so that we are prepared should a risk become a real problem.

From the beginning of a project, the Project Manager and team should have an idea of what the "risk signature" of the project will be. The risk signature will identify expected risks over the course of the project and when the project risks are expected to increase and decrease. During the project, risks should be tracked to determine if mitigation efforts are working.

Risk Management means the entire team is continuously working together to reduce or eliminate risks as part of the normal course of business; not in separate "risk management meetings" that do not involve all team members. Risk Management is not an "add-on;" it must be part of the fabric of project management. As we move forward and continue to open the space frontier, Risk Management must be a part of our project management tool kit. Effective Risk Management depends on a thorough understanding of the concept of risk, the principles of Risk Management, and the establishment of a disciplined Risk Management process. While there is no special set of methods, tools, or communication mechanisms that will work for every project, every NASA manager and employee should have a core set of skills. Centers offer classroom training to bring the team "up to speed" on Risk Management; Web-based training is also available. Let us be serious about Risk Management in order to increase the probability of mission success.

For more information on Risk Management, you may contact the Agency's Office of Safety and Mission Assurance at Headquarters or any Center Safety and Mission Assurance organization.

Fault Tree Analysis

Several investigations are presently addressing the recent failures of the Mars Climate Orbiter and the Mars Polar Lander. NASA is certainly looking forward to the full results of these assessments. Meanwhile, we can improve the potential for success of NASA programs as we await the lessons these teams will report.

I would like to suggest some actions we need to take during the formulation phase of any new program development effort. A few simple steps can increase our chances for preventing failures in our future launches and projects.

In our work, we tend to focus on ways to make things "go right." This confident optimism is an important characteristic that helps us pursue the challenges of invention and exploration. However, to make things "go right," we also need to understand and control the things that can "go wrong." This beneficial pessimism is sometimes a bit more difficult to apply to our own creations, but is needed to increase the likelihood of future successes. Therefore, I ask that we put more effort into analyzing "what can go wrong."

There are a number of engineering tools and techniques that can help us understand the vulnerabilities to our systems. These include the bottom-up analytical approach, known as the Failure Modes and Effects Analysis (FMEA), and the top-down approach, known as the Fault Tree Analysis. A third assessment, the Probabilistic Risk Assessment, integrates information from these two approaches and other sources to assess the potential for failure and help find ways to reduce risk. These analyses constitute a three-pronged approach to help program/project managers assess and mitigate risk and to increase the likelihood for safe and successful missions. This week, I would like to talk in more detail about Fault Tree Analysis.

Fault Tree Analysis is not a new method. The Boeing Corporation first applied it in 1964 to analyze "what could go wrong" with the Minuteman ICBM. It remains, however, one of the best methods for systematically identifying and graphically displaying the many ways something can go wrong. It is "best" in the sense that it is the easiest to use and can be used by anyone, not just safety or reliability experts. It is the easiest in that one begins with a top-level undesired event and works down to identify the subordinate events that could cause such an unwanted outcome. Moreover, in most cases, quantification is not needed to obtain valuable insight into the weaknesses of a design.

At NASA, a Fault Tree Analysis is a methodical review of a system’s hardware and software that begins by envisioning an undesired end state, such as mission failure or loss of crew or vehicle. The project team identifies, in a logical manner, the sequences and combinations of events that could lead to the undesired event. Fault Tree Analysis is most cost-effective when performed early in a project and updated as the project develops. When applied early in the life cycle, it is cheaper to modify a requirement or a drawing than it is to modify hardware or software code later on.

Fault Tree Analysis should also be used to evaluate possible system engineering changes that could eliminate or reduce potential failure paths. As one part of the three-pronged approach, it is a very effective way to find and graphically communicate to engineers and managers a design’s potential "Achilles Heel," should one exist.

Application of Fault Tree Analysis can be beneficial even if initiated late in a program. Questions, doubts, or a late need for additional assurance may sometimes arise. After the Mars Climate Orbiter mishap, a Fault Tree Analysis was done on the Mars Polar Lander, even as it was nearing the end of its long journey to Mars. This analysis was ordered to quickly assess whether the spacecraft might contain latent, but correctable, problems. Ironically, the ability for Fault Tree Analysis to identify what "could" go wrong creates an ancillary capability for helping to find what "did" go wrong after a mishap.

As we prepare for future missions, it is increasingly important that we apply tools such as the Fault Tree Analysis during the formulation and development of a project to ferret out design faults long before any mishap occurs. Think of it as a form of mishap investigation conducted BEFORE there is a mishap.

For our complex and expensive systems, we should not be questioning "why" we should be using Fault Tree Analysis or the other two parts of our three-pronged analysis to ensure our mission’s safety and success. On the contrary, we should be questioning "why not" before we elect to forgo these safeguards. I urge you to understand and employ Fault Tree Analysis to learn "what" can go wrong and to help prevent it when it really counts.

For more information on Fault Tree Analysis, you may contact the Agency’s Office of Safety and Mission Assurance at Headquarters or any Center’s Safety and Mission Assurance organization.

Design for Safety

No one can deny that reaching for the stars is a risky venture, but we should be committed to doing it as safely as possible. We must think safety throughout a program or system life cycle and focus on identifying failure modes and effects in our hardware and software along the way.

A Design for Safety program – a total systems approach addressing safety issues as they are discovered – would help us meet this goal. It would allow us to continuously search for problems and assess risk and solution options from concept through operations. Designing for safety would also include the verification processes to insure the development of safe life designs.

Design for Safety tools would essentially cut the fault tree off at the roots and not allow it to grow. To do this we would depend heavily on learning and knowledge-based tools that will be developed under the Intelligent Systems program. This technology will enable us to create systems that learn and reason for themselves and extract information and knowledge from complex distributed databases. They will also allow us to develop the means to virtually capture the experience and insight of experts. We would use this capability to build high-level "safety oriented" supervisory tools. We would integrate them into the Intelligent Synthesis Environment's life cycle analysis and design tools to develop and institutionalize a smart design process oriented on safety.

How might such a system work?

Design for Safety should start in the concept design phase and continue through the entire life cycle of the project. Design for Safety applies to all project phases – design and development, test and verification, certification, and maintenance and operation. During the design and development phases, Design for Safety tools would conduct automated "what if" studies to evaluate system hazards and their impact on the life and operation of systems. As failure modes are discovered, these tools would quantitatively evaluate safety issues and assess the cost and risk of redundancy versus robustness to minimize risk. And once a system is operational, Design for Safety tools would use the "what if" results and advanced information technology methods to discover patterns and trends and to identify and analyze possible failures throughout the system's life cycle.

They would also track problem reports and maintenance actions to assure our systems were kept in top operating condition. Additionally, operational experience would be used to update analytical models and legacy data/knowledge bases to better predict future system performance and risk. The more experience we gain with our systems, the safer they would be.

We could also use Design for Safety tools and techniques to create a more effective workforce. We could use case studies as educational tools and let people do mock designs under the supervision of a Design for Safety intelligent agent. No tool will replace smart people, but smart tools can create even smarter people and an even stronger NASA.

While this vision requires a long-term commitment to conduct the necessary research and technology development, NASA is prepared to start making it a reality today.